

0.1 Improvements, Cost Reductions & Alternatives

0.1.1 Introduction

For this study, an effort has been made to select specific and feasible technologies giving acceptable performance at reasonable cost. But this is not yet a mature field, and there are many alternative ideas that could be considered. In this chapter we discuss options that might lower the cost, improve performance, or be used as alternatives if unforeseen problems arise. In some cases, cost reductions may be possible without sacrifice of performance; others would hurt performance, but by amounts that would be justifiable by the savings achieved. Some newer technologies might raise performance and lower costs simultaneously.

The discussion is arranged in component order, with the main motive for each modification (Cost, Performance, or Alternative) preceding the subsection titles.

0.1.2 Proton Driver

Performance: Increase protons to $2 \cdot 10^{14}$ (2 MW)

Table 1: Efficiency vs. proton bunch length

<i>rms</i> bunch length (ns)	μ/p	relative
1	0.204	1.02
3	0.20	1.0
6	0.167	0.835
9	0.144	0.72

With an increase of superconducting linac energy (from 1.2 GeV to 1.5 GeV), the proton intensity could be increased by a factor of two (from $1 \cdot 10^{14}$ to $2 \cdot 10^{14}$).

The minimum bunch length achievable is set by the longitudinal emittance of the bunches and by the momentum acceptance of the AGS. For the




Figure 1: Efficiency vs. proton bunch length

baseline 1 MW case, it is hoped to achieve the specified 3 ns rms bunch length without a bunch compressor ring. If the proton bunch intensity is increased by a factor of two, as discussed above, then the bunch length would be expected to increase and the specified rms 3 nsec bunch length could not be achieved without raising the momentum spread above that of the AGS acceptance. Even at the design intensity it is possible that the bunch length would exceed the design value.

The consequences of such an increase in bunch length was simulated, without re optimization. It is not expected that any re optimization will improve the result. The final muons per proton obtained are given in table 1 and figure 1. Note that the cooling system used in this study had larger apertures, and thus higher performance, than the final design, but the sensitivity to bunch length is expected to be the same. It is seen that there is relatively little gain for pulse lengths less than 3 ns (the specified value). For a 6 nsec bunch the efficiency has dropped by 16.5%, and for 9 nsec, the efficiency has dropped by 28%. If such losses of performance are to be avoided, then a bunch compressor ring would be needed.

Performance: Buncher ring

The ring would have fixed superconducting field magnets and could be much smaller in diameter, but with larger momentum acceptance than the AGS. The smaller diameter would reduce space charge effects in the bunched beam, and the larger acceptance allow the short bunches. This ring will have the following features: It will,

1. operate below transition
2. have a small slippage factor, that is, it will be a quasi-isochronous ring
3. have a small dispersion
4. have an acceptance to emittance ratio > 8 , (to be compatible with the tight beam loss limit)
5. have a chromaticity corrector system

Table 2 summarizes the key parameters of the compressor ring.

Table 2: Compressor ring parameters.

Circumference (m)	200
Bending field (T)	4.15
Kinetic energy (GeV)	24
Transition gamma	38.4
η	0.00074
Betatron tune, x/y	14.8/9.2
Maximum beta function, x/y (m)	12.9/19.8
Dispersion function (m)	0.12
Chamber radius (mm)	25
Maximum beam radius, x/y (mm)	7.0/8.6
Acceptance, x/y (m)	48.5/31.6
Beam emittance, x/y (m)	3.8/3.8
Accep./emit. ratio, x/y	12.8/8.3
Natural chromaticity, x/y	-2.5/ - 1.7

In operation, an unmatched bunch is injected from the AGS into the compressor ring. It is extracted immediately after a bunch rotation (bunch

rotation takes a quarter of synchrotron period, i.e. 3 ms, or 4500 turns). Because of the very small slippage, a low rf voltage is required (see Table 3.)

Table 3: RF parameters.

RF frequency (MHz)	5.94
Harmonic number	4
V_{rf} (kV)	200
Bucket height, in dp/p	0.042
Bucket area (eVs)	222
Bunch area (eVs)	10
f_s , center (Hz)	91.5
f_s , edge (Hz)	82.6

The longitudinal parameters of the ring are summarized in Table 4

Table 4: Longitudinal parameters.

	Injection	Extraction
No. particles per bunch (10^{14})	0.17	0.17
RMS bunch length (m/ns)	5/17	0.9/3
Peak current (A)	65	363
Momentum spread (%)	0.4	2.24
Longitudinal emittance (eVs)	10.5	10.5
Broadband impedance ($j\Omega$)	5	5
Space charge impedance $j\Omega$	1.66	1.66
Keil-Schnell threshold ($jM\Omega/m$)	3.75	25.5
Effective rf voltage (kV)	200	248

Clearly, the longitudinal microwave instability threshold will be low at the injection energy, because of the small slippage factor and the low dp/p . To reduce the impedance, the vacuum chamber will have smooth tapered transitions. However, we do not plan to shield the bellows to avoid possible problems with arcing. Despite this we expect to achieve a broad impedance of 5Ω , which is acceptable.

We see from Table 4 that the combination of the broadband and the space-charge impedance is 3.34Ω , slightly lower than the Keil-Schnell (KS)

threshold. Since the beam is below transition, beam instability is not expected. The overall inductive impedance below transition has a focusing effect, which increases the effective rf voltage in the bunch rotation.

Table 5: Transverse parameters.

	Injection	Extraction
Broadband impedance ($jM\Omega/m$)	0.51	0.51
BB imp. induced tune shift	0.0003	0.0017
Space charge inc. tune spread	0.003	0.016
Chromatic tune spread	0.22	1.32
Chromatic frequency (GHz)	59.4	59.4

In Table 5 we summarize the transverse parameters of the compressor ring. We find that the transverse impedance is low, as expected for a small ring ($Z_T \propto R$). Compared to the AGS, the compressor ring is transversely more stable (this is just opposite to the situation in respect to longitudinal instability). The space-charge incoherent tune spread is small and is helped by the strong focusing optics. If the chromaticity is not corrected, the chromatic tune spread is large. This is due to the small slippage factor, the high revolution frequency, and the high tune. For these reasons, we control the normalized chromaticity to about 1%.

The compressor ring design requires very low rf voltage; also, the potential well effect facilitates the short bunch production. The required impedance is reasonable to achieve, and the acceptance/emittance ratio of 8 units is much larger than the one for existing and proposed high intensity proton accelerators; this in conjunction with the large momentum aperture makes it reasonable to expect that the beam loss can be controlled. On the other hand, chromaticity control at the compressor ring is tight, and needs further studies.

Performance: RF for rate to 3.3 Hz (3 MW)

With upgrades of RF and power supplies, the ramp time could be lowered from 150 msec to 100 msec. The cycle time would go from 400 to 300 msec and thus the rate from 2.5 to 3.3 Hz. The cost would be approximately 30 M\$.

Performance: Accumulator for rep rate to 5 Hz (4 MW)

If a full diameter fixed field 24 GeV accumulator was added in the AGS tunnel, then the AGS repetition rate could be increased to 5 Hz. After acceleration, all 6 bunches would be transferred to the accumulator, after which the AGS could immediately ramp down. The bunches would now be extracted from the accumulator at a steady rate of 30 Hz (spacing 33 msec). The cost is estimated at about 50 M\$.

0.1.3 Targets

Alternative: Moving metal band

If unforeseen difficulties made a liquid metal target impossible, then a moving metal band target would be the favored alternative. The performance would be little different from the metal jet. The scheme is discussed in appendix ???. It appears to be feasible at 1 MW, but might not be possible at 4 MW.

Alternative: Carbon

If both the mercury jet and moving band target proved impossible, then a radiation cooled graphite target could be used. This was studied and proposed in Feasibility Study 1[1]. It appears to be a relatively conservative solution (at 1 MW) but would sacrifice a factor of 2 in performance, and require relatively frequent replacement. It is unclear if it could be used at 4 MW.

0.1.4 Capture Solenoid

Cost: Choice of capture field

Figure 2 shows the efficiency for muon production vs. the axial peak field of the capture magnet. Fig 3 gives the magnet cost vs. field. Maximum performance is achieved with the baseline value of 20 T, but the drop in efficiency is small for moderate reductions in this field. A drop from 20 T to 18 T would have an almost insignificant effect (2%), but a 25% saving. A reduction to 15 T gives a 9 % reduction, for a cost saving of 47%.




Figure 2: Efficiency vs. capture field




Figure 3: Magnet Cost vs. capture field




Figure 4: Efficiency of 3 types of inserts in 20 T magnet.

Cost/Performance: Use of wrapped insulation

Figure 4 shows the field vs. power consumption for three different insert coil technologies. The lowest curve is for the baseline design using MgO insulated hollow conductor giving 6T with 12 MW. The dotted line above is for a wrapped ceramic insulation as being developed at MIT[2]. With this conductor, for the same power consumption, the field from the hollow conductor would rise from 6 to 7.6 T, thus lowering the field needed from the superconducting coil from 14 T to 12.4 T, and offering significant savings. Alternatively, the gain in performance could be used to reduce the power consumption.

Cost: Bitter magnet

The upper dashed line in figure 4 is for a Bitter magnet. This technology, which has a very high fraction of its volume available to carry current, is more efficient and cheaper than hollow conductor technology. It is not proposed for a baseline because a suitable radiation resistant insulation would need

development, and even with such insulation it is expected to have a limited lifetime.

Conventional Bitter magnets employ thin wet organic sheet insulation between turns. Ceramic insulation would have to be substituted, but this too would be wet. In the high radiation environment the conductors may corrode. R&D is needed to establish if this is a real concern.

0.1.5 Phase Rotation

Cost: Combining induction linacs 2 and 3

In the baseline design there are 3 induction Linacs. The first linac must be separate from the other two in order to achieve non distorting phase rotation, but the second and third linacs are separate only in order that they each be unipolar. A single second linac with a bipolar pulse approximately equal to the sum of the two opposite polarity pulses would be equally good. This, although slightly less standard appears possible and would be cheaper.

Cost: Fewer induction linacs

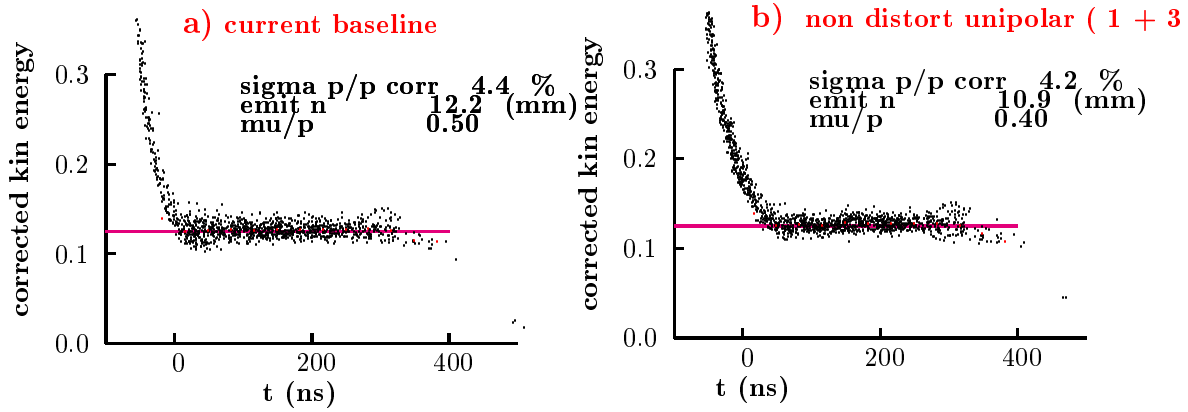


Figure 5: Final energy vs time for different phase rotation systems: a) baseline; b) with IL2 removed.

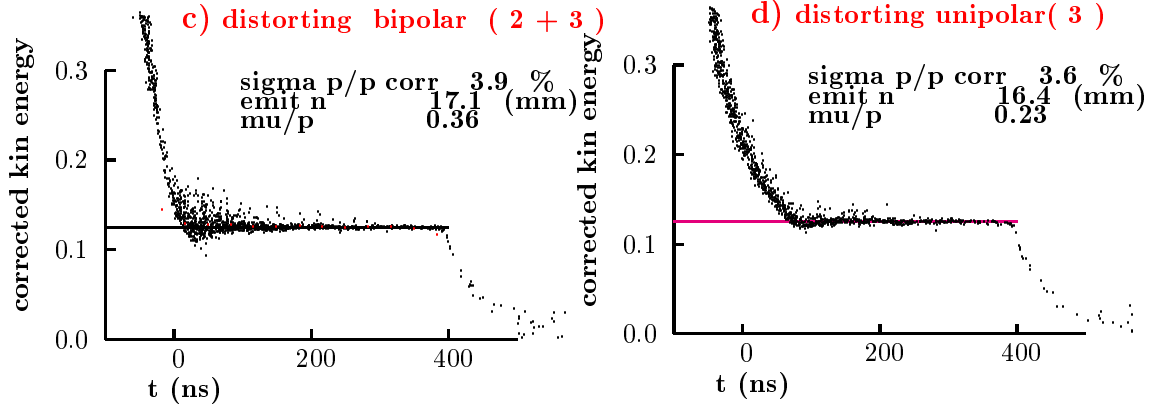



Figure 6: Final energy vs time for different phase rotation systems: c) without IL1; d) without IL1 and IL3.

Further cost savings could be achieved if one or more of the linacs is eliminated and the remaining linacs re optimized. This has been studied assuming a fixed geometrical layout so that the original design could be installed in an upgrade. Figures 5 and 6 show 3 such cases, together with the baseline design. Figure 7 shows the muon production efficiency for the four cases, plotted against the sum of the lengths of the remaining linacs. The losses in efficiency are large when the first linac is eliminated, but less severe if only the second linac is removed (11%).

Performance: Wiggler drifts

The length of the baseline phase rotation is such that there is significant muon loss from decays. A shorter system, using a simple solenoid transport, would result in less time spread and greater final momentum spread, leading to a greater loss in the bunching and cooling sections that follow. A possible improvement could be achieved by using a magnetic chicane or wiggler drift channel that would generate the needed time spread in a shorter distance, and thus less decay loss without sacrificing the final momentum spread.

Figure 7: Efficiency *vs.* length of induction linacs**Performance: Correction of amplitude dependent effects**

In the baseline design, the longitudinal phase area increases by almost a factor of two though the phase rotation. The effect is related to amplitude (it is not present for zero emittance), but is not fully understood. Further study may offer improvements.

Performance: Polarization

A system of double phase rotation has been studied[3] that generated a strong correlation between the muon polarization and final time after the phase rotation. This correlation, though a little diluted, is maintained through to the storage ring and results in correlations between neutrino type and time of detection. The physics need for such correlation has not been well established, and the system requires high gradient (4 MV/m) low frequency (30 MHz) RF close (3-6 m) to the target. The feasibility of such RF has been questioned, but tests at CERN[4] have not ruled it out.

Cost & Performance: Bunched Beam Phase Rotation

Considerable cost savings may be possible by performing the phase rotation with RF after bunching of the un-phase rotated beam[5]. As in induction linac phase rotation, the bunch is first allowed to drift to increase the bunch length and establish a correlation between time and energy, but in this case the bunching is done before the energy is corrected. The RF that performs this bunching is acting on a beam with strong time-momentum correlation; i.e. a beam whose time spread is still increasing with drift distance, and whose sub-bunches, as they are formed, have spacings that are also increasing. This requires that the RF wavelengths used to bunch and hold the bunches, also rise with drift distance. After the bunches have been formed, suitable modifications to the RF frequencies and phases can be employed to accelerate the later bunches and decelerate the early ones, thus ending up with a train of bunches at a single energy, as in the conventional case.

The need to have cavities operating at many different frequencies is a complication. But since the cost of conventional RF acceleration is so much less than that for induction acceleration, the cost of the system is expected to be significantly less. Whether it is as efficient is less clear. For instance, non-distorting phase rotation does not seem possible. But this scheme has the surprising feature of working on both muon signs: the bunches of the opposite sign automatically form between the others. If both signs were subsequently accelerated through the linacs and RLA (injected in the opposite direction), and injected into the storage ring (also in the opposite direction), then a factor of two in efficiency could be achieved. This factor of two might compensate for any lower efficiency in the phase rotation, or it might allow an improved overall performance, while simultaneously lowering the cost.

This solution is far from worked out and many problems remain to be studied. Injection into the ring must be such that timing can be used by the detector to separate the neutrinos from the two different muon trains.

0.1.6 Cooling

Cost: Less cooling

Figure 8 shows the muon production into the defined accelerator acceptance as a function of length. Table 6 shows the values for 4 cooling lengths. Note that the shortest case (47 m) uses only 2.75 m cells. It is seen that a reduction in cooling length from 108 to 88 m, which would offer significant




Figure 8: Efficiency vs. length of cooling

savings, reduces the performance by only 3.4 %, for a saving of ≈ 75 M\$. Shortening the cooling to 47 m lowers the performance by 29 %, a significant loss, but the savings would be large (≈ 230 M\$)

Table 6: Efficiency for three cooling lengths

cooling length (m)	μ/p	loss (%)
108	0.174	0
88	0.168	3.4
68	0.150	13.8
47	0.124	29

Performance: Grid of tubes vs. foil RF windows.

If the radii of the Be foil RF windows could be increased without increasing their thickness, then the performance would be improved. Table 0.1.6 shows results a) for the baseline window thicknesses, and b) for 80 micron Al




Figure 9: μ/p with grids (red upper) vs. foils (blue lower)

Table 7: Performance dependency on RF cavity apertures

Maximum aperture (cm)	μ/p	
	Be foil	Al tubes
21	.174	0.189
25	0.19	0.197
30	0.195	0.205

windows that correspond to a grid of tubes (see below). Note that in each case the length of cooling was adjusted to give maximum performance. In both cases there appear to be significant gains. But for edge cooled Be foils, to avoid excessive temperature rise, their thickness must be increased as the 4th power of radius. If this is done, the performance, instead of rising, falls. For a gas cooled grid of thin walled tubes, however, the pipe thickness is independent of aperture radius and no such problem is encountered.

Tracing with 5 cm diameter pipes has shown that the field non-uniformities

lead to increases in emittance, but these problems can be avoided if the pipe diameters are reduced while their numbers are increased. A second advantage of many small tubes is that, for a given pressure, their wall thicknesses can be reduced.

Given 1 cm diameter pipes, spaced on 2 cm centers, with wall thicknesses of 1 mill (25 μm); then with 1 atmosphere of gas in the pipes, the tension in the walls would be only 3000 psi, which should be acceptable. For such diameters, the non-uniform field effects appear small[6]. When a pair of such grids (at right angles to one another) is simulated by a plane foil with the same average material thickness (80 μm Al), then the performance gain, as seen in table 0.1.6, is 18% for 30 cm apertures and approximately 13% for 25 cm apertures. Fig. 9 shows the accepted muons per proton vs. length for the baseline (lower curve) and with grids and 25 cm radius apertures (upper curve).

Performance: Correlation matching

Within the cooling lattice, particles with high transverse amplitude travel on longer orbits than those on the axis, and thus, for a given momentum, move more slowly in the forward direction. In such a lattice with RF keeping particles within RF buckets, the average forward velocity is controlled by the phase velocity of the RF and constrained to a fixed value. As a result, the stable momenta of particles become dependent on their amplitudes:

$$\frac{dz}{dt} - \beta c \propto A^2$$

where the approximately conserved particle amplitudes:

$$A = \frac{x^2 + y^2}{\sqrt{\beta}} + \frac{1}{\beta} (x'^2 + y'^2)$$

Such a correlation is also generated naturally in the phase rotation process, but, since the phase rotation is done in a different lattice from the cooling, the magnitude of the correlation is not the same. As a result, there is, in the present design, a mismatch in correlation at the entry to the cooling channel. Study is needed to see if performance could be improved by matching these correlations: possibly by raising the solenoid fields used in the transport and phase rotation channels.

Alternative: Limited Flip Cooling Channels

In all solenoid focused channels used for cooling, the axial direction of the field must, at least once, be reversed. If this is not done, canonical angular momentum (i.e. the angular momentum of the beam once out of the axial field) rises and it is impossible to remove it. In the "Super FOFO" lattice, the field is reversed every cell, and significant canonical angular momentum never develops. But there are other solutions with far fewer flips: e.g. single flip[7] or, two flips, as in the example in appendix ???. In these cases the canonical angular momentum is allowed to build up, but is subsequently removed after a flip, by cooling with the opposite field direction. The performance of such alternatives appear to be similar to that of the Super FOFO, but the engineering design of the magnets is very different. The Super FOFO has less stored magnetic energy (often considered an indicator of cost) than the double flip design (about 1/5), but the forces between the coils are higher. At this time, the Super FOFO seems cheaper and at least as efficient, but more detailed engineering will be needed to confirm this choice.

Between the baseline design and the double flip alternative in appendix ??, there is a difference that is not related to their differing lattices. The double flip alternative performs its cooling at higher energy, and an energy that rises along the lattice. This gives a larger longitudinal acceptance, but requires more acceleration for a given cooling. The larger acceptance preserves more muons through the cooling channel, but does not appear to increase the muons accepted by the current acceleration scheme (see sec. 0.1.7).

Performance: Emittance Exchange

In the baseline design, there is a large loss ($\approx 50\%$) of muons as they pass through the cooling channel. This is almost entirely because the lattice does not transport the increases in momentum spread that arise in the absorbers. This loss could be greatly reduced if the longitudinal acceptance of the lattice were increased; for instance by increasing the energies as the beam cools (as done in the double flip alternative of appendix ??). But this would not greatly improve the final result because the longitudinal acceptance is subsequently limited in the acceleration: in particular, in the RLA. Increasing the acceptance there is expensive.

The preferred solution would be to cool the longitudinal emittance. This

can, in principal, be done by introducing dispersion (correlation of momentum with some transverse dimension) and placing a shaped absorber so that the higher momentum particles pass through more material than the lower momentum ones. In such a process, while the longitudinal emittance is reduced, the transverse emittance is increased, and thus we have emittance exchange[8], rather than cooling. But when emittance exchange is combined with transverse cooling, all dimensions would be cooled.

This process, though simple in principle, has been found to be surprisingly hard to achieve efficiently in full 6 D simulations. In unbunched beams, momentum spread can be efficiently exchanged with transverse emittance using bent solenoids or helical magnets. But in such systems, the dispersion introduces time-momentum correlations, which, for bunched beams in the presence of RF, greatly complicate the dynamics. Current designs do achieve some cooling in all 6 dimensions, but with less than ideal efficiency. The problem will, we believe, be solved - it must be solved by the realization of a muon collider - and, when solved, could give up to a factor of two in performance. Further, if the cooled beam emittances could be further reduced, then the needed accelerator acceptances and cost could be reduced.

0.1.7 Acceleration

Performance/Cost: Accelerator acceptance

The acceleration of the muons represents a major cost of the system. This cost could be reduced if the longitudinal and/or transverse acceptances could be reduced. And, conversely, the performance could be improved if these acceptances could be increased. The performance vs. acceptances are given in tb. 0.1.7 and plotted in fig. 10. It is seen that a significant gain in performance could be achieved with greater transverse acceptance, but that the longitudinal aperture accepts almost all muons. Fig. 11 shows a very approximate estimate of the RLA cost vs. longitudinal acceptance[?]. We see that a reduction of the longitudinal acceptance from 150 mm to 100 mm which would reduce performance by 12%, would save 15% of the RLA cost.

Cost: Dog bone configuration

A parametric study of costs[?] has been done on conventional racetrack and dogbone RLA's. The method used a semi-automatic longitudinal motion de-

Table 8: Performance vs. accelerator acceptances

long acc mm	trans acc mm	mu/p	mu/p*
-	-	.231	
150	5	.074	
150	10	.136	
150	15	.174	
150	20	.194	.20
150	25	.205	.216
150	30	.213	.23
50	15	.09	
100	15	.153	
150	15	.174	
200	15	.177	
300	15	.179	

* with shorter cooling to maximize mu/p

sign, minimizing the energy spread. The costs were taken from the First Feasibility study: linac costs proportional to energy gain ($C=38 \Delta E$ per GeV); arc costs proportional to length and energy spread ($C=0.18 \Delta E dp/p$ per GeV and %). The cost units are such that the two RLA's of Study one cost 500 units.

For the conventional racetrack design (fig.12), the method shows that a cost minimum is achieved with 6 turns (fig.14). However, four passes have been chosen for several practical reasons, including the difficulties of designing a switchyard with greater than 4 paths. If these problems could be overcome, then a cost saving of approximately 7% might be achieved.

An alternative geometry for the RLA is the dogbone (fig.13). In this geometry there is only one linac, with the beam passing through it in alternate directions. The minimum cost in this case is found to be at 7 linac passes, is 11% less than the optimized racetrack, and 18% less than the baseline. Despite the larger number of passes, the number of paths on any one side of a switchyard is only four - no more than in the baseline. The savings in a dogbone arise primarily from the ability to reduce the length of the arcs when the momentum is lower. If all the arcs were forced into a single tunnel, the gains would be lost. The above estimates of savings is far from accurate,




Figure 10: Performance vs accelerator acceptance: a) LEFT: transverse (upper line includes re-optimization of cooling length); b) RIGHT: longitudinal.




Figure 11: Relative RLA cost vs approximate longitudinal acceptance

and more study will be needed to determine if the savings are real.

Cost: FFAG

Fixed Field Alternating Gradient (FFAG) acceleration offers the possibility of significant savings. There would be no multiple arcs, and no switchyards:




Figure 12: Schematic of conventional racetrack RLA.



Figure 13: Schematic of dogbone RLA.

the lattice would have a large enough momentum acceptance to circulate the muons from initial to final energy. The number of turns could now be raised, limited only by muon decay considerations, thus lowering the needed RF acceleration per turn.

Lattices have been designed with momentum acceptances over a factor of 2-3. Injection and extraction would be performed using kickers. Designs being studied at KEK[10] employ low frequency, low accelerating gradient RF and accept relatively large decay loss. Work in the US[11] has mainly concentrated on higher gradient superconducting RF with fewer turns and less loss. The main problem in this approach is assuring that the RF phase is right at each pass. The ideal solution is a ring that is exactly isochronous, but the best current designs are less than ideal and require phase control of the RF corresponding to frequency variations of the order of 10^{-4} . This would be easy for conventional RF, but is difficult in a superconducting cavity. The use of ferrites weakly coupled to such cavities is being studied.




Figure 14: Relative costs of RLA's vs. number of passes.

0.2 Summary and Conclusion

Summary

Table 0.2 lists the performance vs. savings and costs for the those options for which numbers are available. Both performance and costs in this table are approximate estimates: far more approximate than those given for the baseline design. They can, however, be used as a guide to what should be studied in the future and how staging might be done.

The cost reductions with no performance penalty are listed in section a) of tb.0.2. If they were all found to be practicable, they would save a total of 98 M\$ (unloaded).

Section b) in tb.0.2 lists the savings and performance penalties of various cost reductions. If they were all implemented, then the cost savings would be 377 M\$, and the performance would be reduced to 54% of the baseline value. An additional cost reduction would be achieved if an initial storage ring for 3 GeV muons were built, and the RLA were initially omitted. The

Table 9: Performance vs. savings in unloaded \$'s

a) COST SAVINGS	%	\$ saved M\$
Wrapped insulation on hollow conductor	0	6
Combined induction IL2 and IL3	0	20
Dogbone RLA	0	72
b) COST REDUCTION	performance loss factor	\$ saved M\$
Capture field	.98	-7
No 2nd Induction linac	.89	-80
60 m less cooling	.71	-230
40 m less cooling	.86	-150
20 m less cooling	.07	-75
100 vs 150 mm long. RLA acceptance	.88	-60
c) UPGRADES	performance gain factor	\$ cost
Increased Linac Energy for $2 \cdot 10^{14}$ ppp	1.72	+12
Buncher Ring	1.16	44
RF & PS's: Rep rate $2.5 \rightarrow 3.3$ Hz	1.32	30
Accumulator: Rep rate $3.3 \rightarrow 5$ Hz	1.52	50
grid + 25 cm apertures + 27 m more cooling	1.13	+100

cost reduction would be substantial (≈ 400 M\$ unloaded) but the event rates far less ($\approx 1/7$).

Section c) in tb.0.2 lists the costs and performance of the upgrades. If all upgrades were done, the additional costs would be 226 M\$ (unloaded) and the performance gain a factor of 4.2.

Optimized Staging

We can now do the following exercise:

1. start with all the cost reductions in place.
2. Consider all possible upgrades and all possibilities of restoring cost reducing items.




Figure 15: Performance vs. unloaded cost for optimized staging.

3. Pick these items in order of their relative performance gain over cost.

This procedure will give the best staging path. Table 0.2 lists the items in order, and fig. 15 plots the performance vs. cost as each item is added.

We note:

- If all cost reductions (except lowering the energy to 3 GeV) are implemented, then the cost savings would be 377 M\$ (unloaded), and the performance reduced to 54% of the baseline value.
- The performance gains from upgrading the proton driver are more "cost effective" than restoring savings from reduced phase rotation or cooling, and should thus be done earlier in a staged approach.
- If the proton bunch intensity is raised to 210^{14} ppp with no bunch compressor, and all other components as for the "minimum cost", then baseline performance would be achieved at a cost 365 M\$ (unloaded) less than the baseline.

Table 10: Performance optimized upgrades

	loss-gain factor	cost	diff/cost %/M\$
Linac $1.2 \rightarrow 1.5$ GeV	1.72	12	6.0
RF & PS's: Rep rate $2.5 \rightarrow 3.3$ Hz	1.32	30	1.1
Accumulator: Rep rate $3.3 \rightarrow 5$ Hz	1.52	50	1.04
Buncher Ring	1.16	44	0.36
Capture field $18 \rightarrow 20$ T	1.02	6.7	0.3
add 20 m cooling	1.21	75	.28
$100 \rightarrow 150$ mm long RLA acceptance	1.12	60	0.2
add 20 m cooling	1.13	75	.17
add 2nd Induction linac	1.12	80	.15
add 20 m cooling	1.03	75	.045
25 cm apertures grids + 27 m more cooling	1.13	100	.13

- For the baseline cost, a performance $3 \times$ baseline is achieved with the fully upgraded (4 MW) driver, with no second induction linac and 20 m less cooling.
- If all upgrades were done, the additional costs would be 226 M\$ (unloaded) and the performance gain a factor of 4.2.

One may also note that upgrading the proton driver to 4 MW would have many other applications:

- neutrino super beam
- rare K and μ decay experiments
- g-2
- spallation neutrons
- etc.

One can thus envisage a staged construction of the project, with no stage costing more than 1 B\$ (with overhead and contingency), and with physics at each stage. For instance:

1. Upgrade driver to 1 MW
2. upgrade driver to 4 MW
3. Minimum cost ν factory at 3 GeV
4. Upgrade to 20 GeV
5. upgrade phase rotation and cooling

Conclusion

Although we believe that the current Study 2 baseline represents a feasible and reasonably costed high performance design, there are many possibilities for cost reduction, performance improvements and staging that may become available, but need further study.

We find that, of the upgrades, increasing the proton driver power to 4 MW appear the most cost effective, and that they should be done in preference to building the full phase rotation and cooling systems.

However, it must be remembered that this section is more speculative than the baseline study, and that caution should be excersized.

Bibliography

- [1] N. Holtkamp, D. Finley, Editors, *A Feasibility Study of a Neutrino Factory Based on a Muon Storage Ring*, Aug., 2000 (http://www.fnal.gov/projects/muon_collider/nufactory/fermi_study_after_april1st/)
- [2] J. Rice, DoE SBIR Phase I Final Report, *Ceramic Insulation for Heavy Ion Fusion and Other High Radiation Magnets*, DoE Grant No. DE-FG03-00ER82979 (2000).
- [3] R.B. Palmer, C. Johnson and E. Keil, *A Cost-Effective Design for a Neutrino Factory*, BNL-66971, CERN SL/99-070 AP, NEUTRINO FACTORY NOTE 09; published in the Proceedings of NuFact99, Lyon, (<http://lyopsr.in2p3.fr/nufact99/>.)
- [4] N. Marseille and W. Pirkel, *Sparking Cavity Test Report*, (<http://nicewww.cern.ch/molat/neutrino/nf52.pdf>)
- [5] D. Neuffer "High Frequency Buncher and $\phi - \delta E$ Rotation for the $\mu^+ - \mu^-$ Source", MUCOOL Note 0181, Oct. 2000, (<http://www-mucool.fnal.gov/notes/>)
- [6] R. Rimmer; grids ***
- [7] V. Balbekov, P. Lebrun, J. Monroe, P. Spentzouris, *The Single Field Flip Cooling Channel for a Neutrino Factory*; MUC Note 0125, March 2000, (<http://www-mucool.fnal.gov/notes/>)
- [8] G. Hanson, (http://needmore.physics.indiana.edu/~gail/emittance_exchange.html)

- [9] S.J. Berg, unpublished
(<http://pubweb.bnl.gov/people/jsberg/talks/010131/0101310all.pdf>);
J.S. Berg, et al., *Acceleration Stages for a Muon Collider*, Proc. PAC99,
New York (1999), pp. 3152.
- [10] Y. Mori, *KEK FFAG Program*, unpublished.
- [11] C. Johnstone, *US FFAG Program*, unpublished.